REMARKS/ARGUMENTS

Claims 1-38 were previously pending in the application. Claims 5, 9, 12, 26, 29, and 36 are amended; and new claims 39-41 are added herein. Assuming the entry of this amendment, claims 1-41 are now pending in the application. The Applicant hereby requests further examination and reconsideration of the application in view of the foregoing amendments and these remarks.

Specification

In paragraph 1 of the office action, the Examiner reminded the Applicant of the proper language and format for an abstract of the disclosure. In response, the Applicant has amended the abstract of the disclosure to have "proper language and format."

Claim Objections

In paragraph 2, the Examiner objected to claims 5, 12, and 29 for certain informalities. In response, the Applicant has amended claims 5, 12, and 29 as suggested by the Examiner. Claim 12 has been further amended to ensure that features (f) and (g) recite steps of a method claim, rather than elements of an apparatus claim. In addition, the Applicant has amended claims 9 and 26 to clarify that the sub-priority of an eligible FTC queue is based on the order in which the eligible FTC queue is set as eligible. The Applicant submits that none of these amendments have been made to overcome any prior-art rejections.

Claim Rejections - 35 USC 112

In paragraph 4, the Examiner rejected claims 10, 27, and 35-36 under 35 U.S.C. 112, second paragraph, as being indefinite.

Regarding claims 10 and 27, the Examiner stated that "it is unclear what a 'sub-queue' is." In response, the Applicant submits that, when a queue is associated with (at least) a high-bandwidth connection and a low-bandwidth connection, the queue can be logically allocated into (at least) two sub-queues, where a higher-priority sub-queue (e.g., corresponding to delay-sensitive traffic) is associated with the high-bandwidth connection, and a lower-priority sub-queue (e.g., corresponding to traffic that is relatively insensitive to delay) is associated with the low-bandwidth connection. See Specification, page 4, lines 19-22, and page 7, lines 14-17.

In response to the rejection of claims 35-36, the Applicant has amended claim 36 to depend from claim 34, rather than from claim 35.

In view of the foregoing, the Applicant submits that the rejections of claims under Section 112, second paragraph, have been overcome.

Claim Rejections - 35 USC 103

In paragraph 6, the Examiner rejected claims 1-4, 8, 10, 12-13, 15-21, 25, 27, 29-30, and 32-38 under 35 U.S.C. 103(a) as being unpatentable over Fan in view of Rusu. In paragraph 7, the Examiner objected to claims 5-7, 9, 11, 14, 22-24, 26, 28, and 31 as being dependent upon a rejected base claim, but indicated that those claims would be allowable if rewritten in independent form. For the following reasons, the Applicant submits that all of the now-pending claims are allowable over the cited references.

Claims 10 and 27

Claims 10 and 27 depend from claims 9 and 26, respectively. Since the Examiner indicated that claims 9 and 26 were directed to allowable subject matter, the Applicant submits that claims 10 and 27 should also be directed to allowable subject matter. The Applicant requests confirmation that claims 10 and 27 are directed to allowable subject matter.

Claims 1, 18, and 38

According to claim 1, an occupied FTC queue provisioned for burst scheduling is identified as a super-occupied FTC queue when the number of cells enqueued is greater than a specified number. Claim 1 also recites that, when the super-occupied queue is serviced, the number of cells dequeued is based on a burst size.

In rejecting claim 1, the Examiner admitted on page 4 that "Fan et al do not disclose (b) identifying an occupied FTC queue provisioned for burst scheduling as a super-occupied FTC queue when the number of cells enqueued is greater than a specified number; and when the super-occupied queue is serviced, the number of cells dequeued is based on a burst size."

The Examiner cited Rusu as providing the features missing from Fan. In particular, the Examiner stated on page 5 that "Rusu et al disclose in Figure 2 a switch system wherein the circuitry accelerates the output transfers from queues 101 when the queues 101 are close to filling up (greater than a specified number), as measured by a differential queue length mechanism," citing column 5, lines 52-60. The Examiner stated further that "When the queue 101 that is close to filing up is serviced, the number of cells dequeued is based on a burst size (maximum number of cells that may be subtracted from the queue in the programmed time interval)," citing column 6, lines 20-27. For the following reasons, the Applicant submits that the Examiner mischaracterized the teaching in Rusu in rejecting claim 1.

Rusu teaches a technique for keeping track of the number of cells stored in each of a number of different queues, which number Rusu refers to as the "queue length." As indicated in Fig. 2, each queue has a differential measurement counter 120, which is used to keep track of the number of cells stored in the corresponding queue. See column 5, lines 3 *et seq*. Counter 120 includes current queue length counter 202, which is incremented every time a cell is written into the corresponding queue and decremented every time a cell is read from the corresponding queue. See column 6, lines 4-6. At fixed intervals, equivalent to n cells time (see, e.g., Abstract, lines 7-10), the value in counter 202 is compared to the previous queue length 201 by ALU 203 to generate a differential queue length, which is equal to the change in the queue length from the last measurement. See, e.g., column 6, lines 31-38. These differential queue length measurements can be used to control the speed of cell flow in or out of the corresponding queue. See column 8, lines 48-51.

Claim 1 is directed to a method of selecting queues for service. In Rusu, queue server 205 of Fig. 2 controls the selection of queues for service (i.e., determining the order in which queues are selected for extraction of data cells stored in those queues). The only teaching in Rusu directed to the selection of queues for service is found in column 6, lines 7-19. According to these very limited teachings, "There are many ways for the queue server (205) to provide output of the individual queue data, including interleaved based on a fixed cycling to a clock, round robin, prioritization based on a priority register as to which queue is to go in which order, etc." Rusu provides no other teachings related to the selection of queues for service.

The three scheduling techniques identified by Rusu (i.e., interleaved based on a fixed cycling to a clock, interleaved based on round robin, and interleaved based on prioritization based on a priority register as to which queue is to go in which order) are all conventional scheduling techniques that determine the order in which queues are selected for service. Significantly, in each of these conventional scheduling techniques, when a queue is selected for service, (at most) only a single cell is extracted from that queue. Depending on the details of the scheduling technique and/or the particular conditions of the queues, it may very well be that the scheduling technique will select the same queue two or more consecutive times (for example, when only one queue has any cells enqueued). Nevertheless, at each of these different queue selections, no more that one cell is ever extracted from the selected queue.

In the first passage cited by the Examiner in rejecting claim 1 (i.e., column 5, lines 52-60), Rusu teaches that "Once the rest of the switching circuitry knows when the queues are close to filling up, the connected circuitry can slow down its input transfers to the queues or accelerate its output transfers from the queue, or make other decisions." Depending on the particular features of the corresponding scheduling algorithm, there are different ways in which a conventional scheduling technique (i.e., a scheduling technique in which at most one cell is extracted at each selection of a queue) can "accelerate its output transfers" from a queue. For example, a priority-based scheduling technique can accelerate the output transfer from a particular queue by assigning that queue a higher priority level. In that case, the scheduling technique will select that queue more frequently, thereby effectively accelerating the rate of data output from that queue. Note that increasing the priority level does not change the fact that at most one cell is extracted from the queue every time the queue is selected for service.

In the second passage cited by the Examiner in rejecting claim 1 (i.e., column 6, lines 20-27), Rusu teaches that the size of the register used to represent the differential queue length should be large enough (i.e., have enough range) to prevent counter overflow or underflow. This has absolutely nothing to do with the number of cells that are removed from a queue each time the queue is selected for service. As such, this passage does not provide any support for the Examiner's conclusion that Rusu teaches "When the queue 101 that is close to filing up is serviced, the number of cells dequeued is based on a burst size (maximum number of cells that may be subtracted from the queue in the programmed time interval)." The Applicant submits that this statement constitutes a misrepresentation of the teachings of Rusu. In particular, Rusu provides no such teachings or even suggestions, either in the passage cited by the Examiner or anywhere else in that document.

In conclusion, the Applicant submits that neither Fan nor Rusu provide any teaching or even suggestion (1) that an occupied FTC queue provisioned for burst scheduling is identified as a super-occupied FTC queue when the number of cells enqueued is greater than a specified number or (2) that, when the super-occupied queue is serviced, the number of cells dequeued is based on a burst size.

For all these reasons, the Applicant submits that claim 1 is allowable over the cited references. For similar reasons, the Applicant submits that claims 18 and 38 are allowable over the cited references. Since the rest of the claims depend variously from claims 1, 18, and 38, it is further submitted that those claims are also allowable over the cited references.

New Claims 39-41

To the extent that there may be any ambiguity in the recitations of claims 1, 18, and 38, new claims 39-41 explicitly state that more than one cell is dequeued from the super-occupied queue during a single selection of the super-occupied queue for service. (Note that it might be possible for a super-occupied queue to have one or no cells currently enqueued. In those particular situations, a single

selection of that queue for service would obviously not involve the dequeuing of more than one cell. To account for such situations, claims 39-41 recite "a single selection," rather than "each single selection.")

In view of the foregoing, the Applicant submits that the rejections of claims under Section 103(a) have been overcome.

In view of the above amendments and remarks, the Applicant believes that the now-pending claims are in condition for allowance. Therefore, the Applicant believes that the entire application is now in condition for allowance, and early and favorable action is respectfully solicited.

Respectfully submitted,

Date: 12/19/05

Customer No. 46900
Mendelsohn & Associates, P.C.
1500 John F. Kennedy Blvd., Suite 405
Philadelphia, Pennsylvania 19102

Steve Mendelsohn
Registration No. 35,951
Attorney for Applicant
(215) 557-6657 (phone)
(215) 557-8477 (fax)